FLASHARRAY™ AND THE GDPR

How Pure Storage® FlashArrays can help enterprises comply with the data processing and storage provisions of the European Union's General Data Protection Regulation.

EXECUTIVE SUMMARY

The European Union's General Data Protection Regulation (GDPR) went into effect on May 25, 2018. The regulation defines handling, use, and transfer requirements for entities that deal with the personal data of EU residents, as well as those individuals' rights with respect to their data. GDPR applies to all entities that handle the personal data of EU residents, regardless of whether the entities are in an EU member country. The regulation applies to the processing and storage of data in digital form, regardless of where processing and storage occur.

THE SCOPE OF THE GDPR

The GDPR declares that EU residents (called *natural persons* and *data subjects* in the regulation) own their personal data, and specifies both their rights with regard to it, and obligations of entities that acquire and process it, particularly with respect to keeping it secure and available.

Individuals' rights to their personal data include:

- The right to access it
- The right to rectify errors in it
- The right to know how it is being processed and to restrict the types of processing it undergoes (within certain legal limits)
- The often-cited *right to be forgotten* (i.e., to have personal data destroyed when it is no longer required for legitimate purposes).

The regulation classifies entities that deal with individuals' personal data as either:

Controllers
Entities that determine the purposes and means of processing personal data, or

Processors
Entities that perform processing tasks as instructed by controllers.

> ***Article 1***
>
> **Subject-matter and objectives**
>
> 1. This Regulation lays down rules relating to the protection of natural persons with regard to the processing of personal data and rules relating to the free movement of personal data.
>
> 2. This Regulation protects fundamental rights and freedoms of natural persons and in particular their right to the protection of personal data.
>
> 3. The free movement of personal data within the Union shall be neither restricted nor prohibited for reasons connected with the protection of natural persons with regard to the processing of personal data.
>
> *Excerpt from GDPR Article 1 defining the regulation's purpose and scope*

An entity can fulfill both roles. In the context of the regulation, the term *processing* encompasses both:

Manual operations
: Acquisition, filing, alteration, and disclosure, etc.

Automated operations
: Electronic processing, storage, and transmission of data in digital form.

The GDPR regulates controllers and processors as to the types of personal data they may acquire and the purposes for which they may process it, and specifies protections they must provide against both unauthorized access and loss or destruction during processing, storage, and transfer. Additionally, it obliges processors to disclose what personal data they store and process to its owners, to rectify verifiable errors in it, and to destroy it when it is no longer relevant to its intended purposes. Finally, the GDPR specifies procedural mechanisms for compliance and lays out penalties for non-compliance.

Thus, the GDPR deals with both

Policy
: What data may be collected and for what purposes it may be used, the rights of EU residents with respect to their data, etc.

Technology
: How data in digital form should be secured against unauthorized access and destruction as it is processed, transferred, and stored.

COMPLYING WITH THE GDPR

Since May 25, 2018, entities that acquire and process the personal data of EU residents have been required to comply with the GDPR. Compliance requires that controllers and processors use verifiable procedures to prevent *personal data breaches*, defined in the regulation to be events that lead to "*the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data transmitted, stored or otherwise processed*."[1]

In terms of policy, compliance encompasses organizational structures, data handling procedures, physical security of data repositories and processing facilities, and hiring, training, and auditing operations. From a technology standpoint, compliance consists of ensuring that computing hardware, software, storage, and communication facilities, when properly managed and maintained, provide high barriers to theft, unauthorized disclosure, alteration, and inadvertent or malicious destruction of EU residents' personal data.

[1] Official Journal of the European Union, 4.5.2016, Article 4(12).

A MODEL FOR HANDLING DIGITAL DATA

Figure 1 is a simple model of a lifecycle for personal data in digital form. The numbers in the figure represent points at which personal data must be explicitly secured.

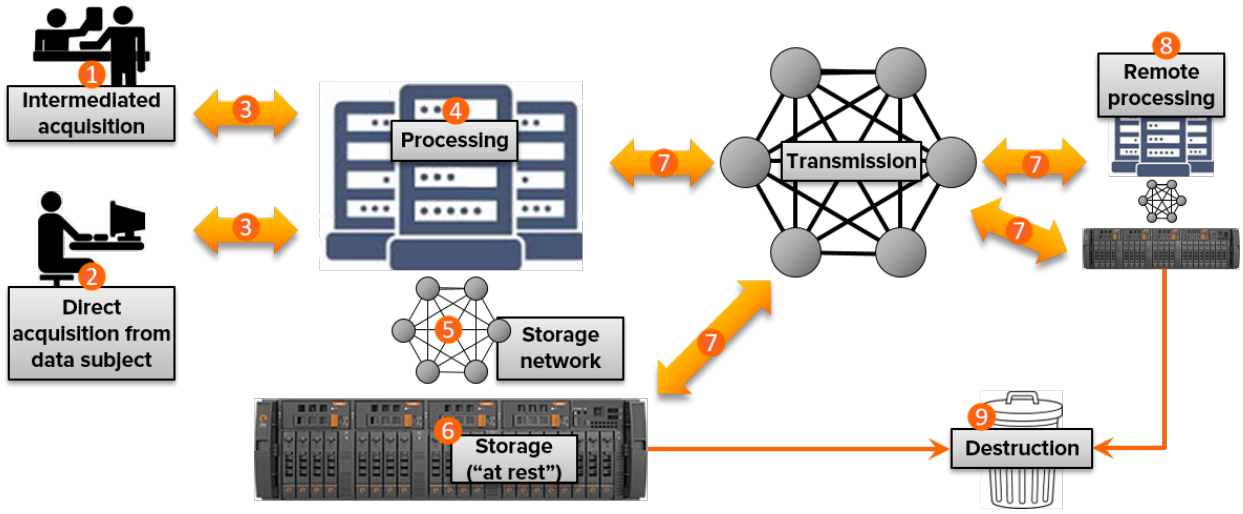

Figure 1: Digital Data Acquisition, Processing, Transmission, and Storage

1. Some digital data originates when a data subject interacts with an intermediary, e. g., a bank teller, salesclerk, or government official. The intermediary interacts with a computer system on the subject's behalf. Privacy and security depend on the trustworthiness of subjects, intermediaries, and processing systems, e.g., via robust vetting and authentication.
2. Increasingly, data subjects digitize their own personal data by interacting with ATMs, governmental and private sector websites, and so forth. Responsibility for securing personal data lies primarily with the networks and computer systems that the subjects use to interact.
3. Some communication links between data subjects, processors' agents, and processing facilities are permanent, and have strong security. But increasingly, data subjects and processors' agents use wireless (e.g., mobile credit card readers) and semi-public Internet access points to interact with processing facilities. These links must be secured against unauthorized access and passive "snooping."
4. Most processing of personal data subject to the GDPR takes place in physically secure data centers. Security of data during processing requires trustworthy, well-managed and maintained hardware and software, comprehensive auditing, and strict control of human access to computer systems.

5 The majority of data processors subject to the GDPR use storage networks to connect their servers to storage. For storage networks contained within a data center, data is secured while in transit by restricting access to facilities and equipment. For networks with external connections (e.g., to servers and storage in separate facilities or to remote replication targets), data must be secured by network-encryption while in transit.

6 Storage systems that hold data "at rest" must protect it not only against unauthorized access and alteration, but also against failure and theft of the systems or their components. Storage systems should encrypt stored data, manage encryption keys securely, closely control human access, and indelibly log all administrative actions.

7 Processors transmit both individual data items (e.g., credit/debit transactions) and bulk data (e.g., for analysis, testing, or archiving). Ideally, private networks would be used, but most processors use common carrier facilities. Typically, they secure transmission over public networks with virtual private networks (VPNs), encryption, or a combination of the two.

8 Sending personal data to a remote system creates a *copy*. The GDPR declares that data subjects have the right to know what copies of their personal data exist and the purposes for which they are used. In addition, they have a "right to be forgotten"—for their personal data to be eradicated once there is no longer a legitimate reason to retain it. To eradicate data, one must know where it is, so data processors are obliged to track copies as they move among their own and their processing partners' facilities.

9 The right to be forgotten requires that processors eradicate personal data upon owner request when it no longer serves a legitimate purpose. For individual items, responsibility for complying with this provision lies with the processor's procedures for dealing with data subjects. For bulk destruction of data sets, such as survey results, some storage systems support fast reliable eradication of large blocks of data.

PURE'S ROLE IN GDPR COMPLIANCE

As a supplier of data storage systems, Pure's key contributions to GDPR compliance are:

- Protecting data both "at rest" in its storage systems and in transit during replication between them (points 6 and 7)
- Keeping data available and secure as it moves between application servers and storage (points 5 and 8) over storage networks with connections outside the data center. This typically requires coordination with storage network and/or server facilities.
- Ensuring to the extent possible that no personal data is exposed in the logs that systems transmit to the company's Pure1 Virtual Private Cloud for analysis and troubleshooting. Technical Brief TB-210301[2] describes the role of Pure1 in GDPR compliance.

The remainder of this brief illustrates how FlashArray storage can be a key component of an organization's comprehensive GDPR compliance strategy.

[2] https://support.purestorage.com/Pure1/Pure1_Security/Pure1_Security_Reference/Pure1_and_GDPR_Compliance

FLASHARRAY AND GDPR COMPLIANCE

As a supplier of digital data storage systems, Pure is committed to keeping the data stored in its systems both available to authorized users and secure against electronic intrusion[3] and physical misappropriation.

Availability and security measures implemented by FlashArray include:

Data Availability

Arrays keep data intact and accessible during all single-component failures and many concurrent failures of multiple components

Access Control

Arrays restrict administrative access to credentialed individuals, each associated with a role that only permits specific actions. There are no mechanisms for administrators to access or modify stored data

Data Protection

- Arrays encrypt all staged and stored data and metadata using the well-known AES-256 algorithm. Encryption cannot be disabled
- Optional *Key Management Interoperability Protocol* (KMIP) servers and removable SmartCards protect data in situations where physical security of arrays is problematic

Data Integrity

Immutable snapshots provide unalterable records of data set contents at key points in time.

Together, these properties help data controllers and processors "design GDPR compliance by default" as they implement new processing systems.[4] Technical Briefs TB-160201, *Securing FlashArray "Data at Rest"* and TR-160202, *FlashArray Access Security*[5] describe how FlashArrays protect stored data from loss and unauthorized access, even under adverse conditions. When combined with strong network security and robust system access and data handling policies on the part of processors and controllers, FlashArrays are an important component of an overall GDPR compliance strategy that is both comprehensive and cost-effective.

---

[3] Pure executes Data Protection Agreements (DPAs) with its vendors who process personal data related to prospective and existing customers in fulfillment of service contracts with such customers. Pure uses intra-company Standard Contractual Clauses to fulfill requirements for transfer of personal data from the European Economic Area (EEA) to the United States in accordance with the GDPR.

[4] Excerpt from GDPR Article 25 (Data protection by design and by default): "…the controller shall, both at the time of the determination of the means for processing and at the time of the processing itself, implement appropriate technical and organisational measures, such as pseudonymisation, which are designed to implement data-protection principles, such as data minimisation, in an effective manner and to integrate the necessary safeguards into the processing in order to meet the requirements of this Regulation and protect the rights of data subjects."

[5] https://support.purestorage.com/FlashArray/FlashArray_Security/FlashArray_Security_Reference/The_FlashArray_Data_Security_Model_TB-160201
https://support.purestorage.com/FlashArray/FlashArray_Security/FlashArray_Security_Reference/The_FlashArray_Data_Security_Model_TB-160201

FLASHARRAY AND COMPLIANCE: KEEPING DIGITAL DATA AVAILABLE

FlashArray provides robust facilities for keeping data intact and accessible. The arrays are designed to continue operating in the presence of hardware component failures, including failure of an entire controller. RAID-HA protection makes data recoverable from all single and concurrent double read failures, as well as many failures that affect more than two devices. Intra-device checksums detect *latent* errors that deliver corrupt data or data from the wrong locations.

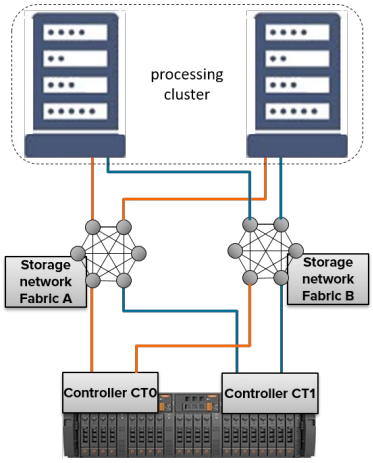

Figure 2: Redundant Sever to Storage Connections

Arrays also protect against administrative errors by imposing an automatic 24-hour delay when volumes are destroyed. During the delay volumes destroyed in error can be recovered.

FlashArray provides these protections along with data reduction that reduces raw storage cost with virtually no impact on I/O performance. Technical Report TR-150302[6] describes how the arrays protect stored data against loss due to component failure and unintended destruction.

Arrays keep data available to application servers ("hosts") via storage network ports on both controllers. When arrays are cross-connected to clustered hosts via two independent storage network fabrics (Figure 2), data remains accessible for processing if a storage network path, a host, or an array controller should fail.

Similarly, processors can configure connections between replicating FlashArrays redundantly so that replication survives network and controller port outages (Figure 3).

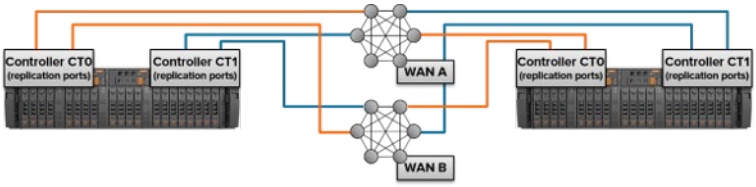

Figure 3: Redundant Replication Network Paths

[6] https://support.purestorage.com/FlashArray/PurityFA/FlashArray_Technical_Reports/FlashArray_Technical_Papers_-_NDA_Distributed/FlashArray_Data_Protection%3A_RAID-HA_NDA_TR-150302
(requires non-disclosure agreement).

FLASHARRAY AND COMPLIANCE: KEEPING DIGITAL DATA SECURE

Digital data storage does not exist in a vacuum. Its inherent purpose is to create a durable record of manipulations of data by data processors or by subjects. Storage is integral to the digital data lifecycle represented in Figure 1, both while data is "at rest" and while it is "in transit" between processing systems and arrays.

For example, storage networks that connect application servers to storage (point 4 in Figure 1) carry personal data. For "air gapped" storage networks entirely within a data center, this data is typically protected by data processor policies that limit access to storage networks and to the servers and storage systems connected to them. Processors should restrict system access to specific trusted individuals, and specify administrative roles narrowly, with disjoint responsibilities. Different individuals should manage storage, servers, networks, and security.

Data on storage networks that have connections outside the data center should be protected by combinations of VPNs and encrypting network switches.

Systems should log all administrator actions. To comply with the GDPR, data processors must use IT products that implement policies to protect personal data and enforce those policies.

FlashArrays contribute to GDPR compliance by securing stored data against theft, either electronic (e.g., unauthorized access by an administrator or host computer), or physical (e.g., misappropriation of storage devices or entire arrays). But in addition to the arrays, security in other IT components as well as processors' policies are required for full compliance.

The key FlashArray features that raise barriers to unauthorized data access are:

No administrator access to stored data

The arrays' command line, graphical, and REST administrative interfaces do not include facilities to allow stored data to be retrieved or altered by any administrator role.

Controlled administrative access

Administrator accounts are both credentialed and role-based. Arrays can manage credentials locally or integrate with Active Directory services for centralized access control. Administrators' activities are limited to their roles—monitor-only, storage management, or full access—regardless of which interface is used to gain access.

Controlled access by host computers

FlashArrays only execute I/O commands from host computers that have been explicitly *connected* to specific *volumes* (virtual storage devices analogous to conventional disks) by an authenticated administrator in the storage or the array administration role.

Audit logging

Arrays log all administrator actions. They store logs locally and can also be configured to forward them to syslog servers. Arrays connected to the Internet, as most are, transmit logs to the Pure1® Cloud, where they can be viewed by Pure's Technical Support Engineers (TSEs) or audited by agents of the array owner through Pure1 Management Services.

Encryption of all stored data all the time

FlashArrays encrypt all stored data and metadata using the widely accepted AES-256 algorithm.[7] Arrays unlock their flash devices at power-on, using device-specific *access keys* regenerated at least every 24 hours. They encrypt their encryption and access keys, and store partitions of them in their devices, so that more than half an array's devices must be present for key regeneration. For additional protection where physical security is a concern, arrays can be integrated with remote KMIP servers or use removable *smartcard* readers. When either of these is in use, keys cannot be regenerated without access to them.

Immutable snapshots

FlashArray snapshots cannot be altered. Thus, they protect against inadvertent destruction of data, and provide point in time records that can be used to detect unauthorized changes or to recover from ransomware and other malware attacks.

FLASHARRAY PHYSICAL SECURITY

Data processor policies frequently centralize the management of all encryption keys, regardless of the keys' purposes or the locations at which they are used. FlashArray manages data encryption keys autonomously, but for additional security, arrays can be configured to integrate with Data Security Management (DSM) servers that use the *Key Management Interoperability Protocol* (KMIP). When KMIP is in use, FlashArray and DSM server exchange cryptographically signed certificates to establish a trust relationship, after which the DSM server helps the array to decrypt its key (which is stored on the array encrypted). Without a connection to the DSM server, an array cannot recover its key, and so can neither read nor store data.

Where an array's physical security (for example, against device theft) cannot be guaranteed, optional removable *smartcards* can be installed in an array's controllers. The smartcards contain tokens that controllers use to unlock flash devices and to recover the array's data encryption key. If the cards are removed and the array power cycled, data can neither be read nor written.

Both KMIP and smartcard enablement are one-way operations. Once either is engaged, disengagement makes all data on the array permanently inaccessible.

[7] The United States Federal Government certifies that the arrays comply with the FIPS 197 and FIPS 140-2 Level 1 data security standards. See https://csrc.nist.gov/projects/cryptographic-module-validation-program/Certificate/2467

FLASHARRAY AND THE "RIGHT TO BE FORGOTTEN"

Several articles of the GDPR relate to a data subject's "right to be forgotten"—the right to request that data controllers and processors destroy data that no longer serves a legitimate purpose. This is actually two requirements:

Destruction of a single subject's data items

Deleting a subject's records requires structural knowledge of the files or databases that contain them and the applications that process them. Thus, responsibility for compliance lies with users of applications that process files or databases and with processor data handling procedures. It must be possible to identify, locate, and remove all instances of a subject's records from data sets that contain them (e.g., snapshots, copies, backups, and other transformations).

Application and database "deletion" typically makes records inaccessible, but does not physically obliterate them immediately. Data processors must physically destroy all copies of a subject's personal data within an acceptable time after they are deleted from an application retrieval point of view.

Destruction of data about entire groups of subjects

Bulk personal data, for example obsolete polling results or database tables, can often be destroyed somewhat differently. Deleting hundreds of thousands of records one by one is usually too onerous to be practical. Most storage systems, including FlashArray, support an *unmapping* function that overwrites ranges of blocks with zeros upon server command. This can expedite destruction of large numbers of personal records, but again, data processor policies must guarantee that all instances of such data sets are destroyed.

FlashArrays present storage to application and database servers as disk-like virtual *volumes*. The arrays have no awareness of volumes' structures or contents. While a FlashArray administrator with the storage or array role can destroy an entire volume on command,[8] s/he cannot alter or erase individual data items in it. Thus, destroying a FlashArray volume to erase large quantities of data is only possible when the data to be erased is the sole occupant of the volume. This is of particular concern with older VMware implementations that create multiple files on a single FlashArray volume, each one representing a virtual machine's system disk. In such scenarios, VMware facilities must be used to erase these "vmdk" objects, which a FlashArray perceives as block ranges on a single volume.

[8] *Destroying* a FlashArray volume has a 24 hour delay that can be truncated by a command to *eradicate* its contents immediately.

THE DATA ENCRYPTION DILEMMA

Data processors generally accept that encryption of personal data throughout its digital lifecycle will eventually be a practical necessity for GDPR compliance. Figure 1 suggests that "end to end" encryption can occur in multiple stages. For example, data should be encrypted on the communication path between origin and processor (point 3), but it must be decrypted for processing (point 4).

As a practical necessity, many storage networks cannot be contained within a data center, particularly Ethernet-based ones. But even within a data center, encrypting data at its source limits the number of points at which it is "in the clear" while it moves or is stored. Some consultants and equipment vendors recommend that sensitive data be encrypted at the application server and transferred and stored in encrypted form. Data encrypted at the server is protected—in primary storage, when copied to other servers, and when backed up or archived. But it requires (a) managing encryption keys so that copies can be decrypted for legitimate purposes, and (b) enforcing policies that ensure that only authorized individuals and systems have access to keys and data, and only for legitimate purposes.

Data encryption at the application server covers many GDPR compliance requirements, but it carries a high cost. *Data reduction*—the elimination of redundancy prior to storing data persistently—is a common expectation of data processors as they evaluate storage alternatives. Increasingly, processors budget for and acquire storage based on estimates of data reducibility rather than on raw capacity requirements. Reducibility varies with the type of data, but often falls in the range of 3:1—5:1. Because so much personal data consists of character strings, it typically reduces at the higher end of this range. Put another way, with reduction, processors can reduce physical storage requirements by 67-80%. Acquisition cost, space requirements, and power consumption are proportionally lower as well.

Reduction removes redundancy from data in two ways:

- *Compression* replaces sequences of bytes repeated within a block with more compact representations
- *Deduplication* replaces entire blocks that are identical to already-stored blocks with pointers to a single instance.

Encryption inherently produces quasi-random bit patterns which essentially eliminates the possibility of compression, and severely limits possibilities for deduplication.[9] Thus, even ignoring the computational

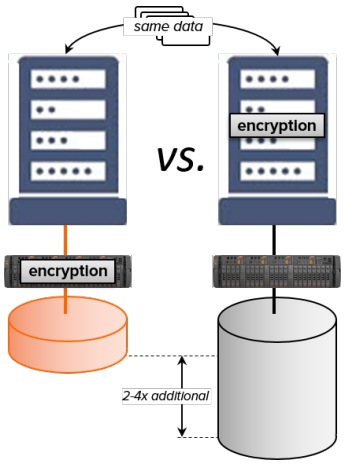

Figure 4: The Cost of Encryption at the Application Server

[9] In principle, identically aligned blocks of identical data could encrypt to identical bit patterns, and therefore, despite such encryption being weak, would be amenable to deduplication. One factor that makes FlashArray data reduction

impact of encryption on applications, encrypting data at the server eliminates the cost advantage of data reduction that most data processors depend on.

Server-side data encryption is only the tip of the storage cost iceberg. Each time data is copied—for analytics, for development testing, for backup, or for disaster protection—the lack of reduction means that three to five times as much storage is consumed compared to data encrypted by storage systems.

Moreover, the keys used to encrypt data at the server are an inherent security weakness. For example, the key that encrypts a production data set must be available to systems that analyze copies of it, use copies for development testing, restore backups, and so forth. Every use of the data set widens the circle of systems and individuals with access to its contents. Moreover, if new versions of a production data set are encrypted with different keys, all users of copies must track keys so they can use the correct one for each version of the data set they process.

Server-side encryption is thus a "brute force" solution to one area of GDPR compliance. Assuming that key proliferation can be managed, it does secure data, but at significant cost to the processor by sacrificing one of the most important storage technology advances of the past decade—data reduction. The alternative, encrypting in the storage system, as FlashArray does, preserves the cost advantage of reduction and mitigates the key management problem, but unless data is encrypted by the storage network, it is exposed as it moves between servers and storage systems, or as it is replicated between storage systems.

A COMPROMISE: SELECTIVE PSEUDONYMIZATION[10]

Application servers can encrypt data at different levels. They may encrypt every block sent to a storage system, data written to a specific file system or database, or to specific files or database columns (e. g., individuals names or identifying characteristics). Because GDPR is specific about the types of personal data that are subject to protection, data processors can comply by adopting server-side encryption or other types of pseudonymization only for sensitive items, with other data stored in the clear. The intent would be to comply with GDPR digital data protection provisions while retaining at least *some* of the cost benefit of data reduction. Selective encryption is likely to lessen the degree to which a storage system can reduce data, but where the amount of non-sensitive data is significant, some benefit should still accrue. Moreover, selective pseudonymization is likely to consume less server processing power than bulk encryption of all data processed and stored.

more effective than other schemes, however, is its ability to deduplicate blocks with different sector alignments. Encryption at the application server would eliminate this advantage.

[10] The term pseudonymization is used in the text of the GDPR regulation.

PROTECTING REMOTE DATA

For storage networks contained within a secure data center, data is protected by a combination of (a) controlled access to servers, storage networks, and storage systems, (b) separating and narrowing administrative roles, (c) auditing administrative actions, and (d) scrupulously maintaining application and environmental software and firmware.

But these measures do not protect data in transit to and from remote servers and storage systems. Data processors may find it simpler to comply with the GDPR by encrypting all data in transit, regardless of network technology and topology.

Network vendors offer hardware and software encryption tools for TCP/IP and Fibre Channel networks, so while GDPR compliance may require some network architecture redesign, it is possible to provide robust protection for data throughout its digital lifetime along with the cost benefit of data reduction.

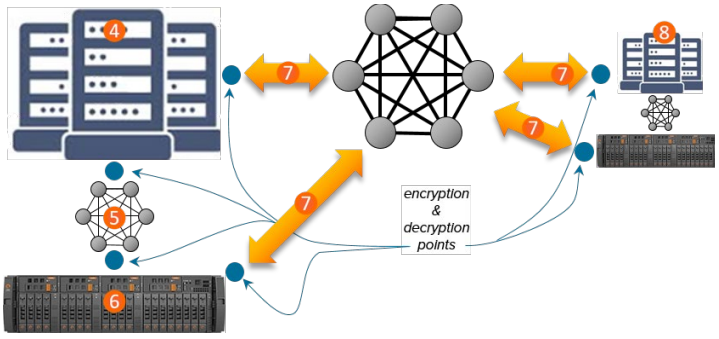

Figure 5: Encrypting Data in Transit

Encrypting data in transit provides reasonable protection against “snooping” and “man-in the middle” attacks, but data remains susceptible to *threats from within*—administrators that divulge keys indiscriminately, connect storage arrays to unauthorized servers, and so forth. Thus, even end-to-end encryption of data must be accompanied by policies with interlocking safeguards against misappropriation and use by the data processor’s own personnel.

THE BOTTOM LINE

- Starting in May 2018, GDPR compliance is required to do business in the EU.
- GDPR contains both organizational, procedural, and digital data handling provisions. Digital data handling includes both keeping personal data available for legitimate use and protecting it from misappropriation and misuse.
- GDPR does not specifically require encryption of digital data, but end-to-end encryption simplifies compliance to the point where many data controllers insist upon it, and data processors tend to adopt it as a default practice.
- Encryption at the application or database server protects data throughout its digital life, but at a substantial storage cost due to the near impossibility of reducing encrypted data. The cost is multiplied each time encrypted data is copied for analytics, backup, or testing.
- Server-level encryption makes key management complex and introduces inherent security weaknesses due to the need to distribute keys to every user of a data set.
- For systems like FlashArray, that both reduce data and encrypt it for storage, end-to-end encryption that preserves both the cost advantage of reduction and operational simplicity can be achieved with encrypted network links between servers and storage and between replicating pairs of arrays.
- FlashArrays control access to administrative operations, log every administrator interaction, and encrypt all stored data and metadata *all the time*—encryption is not selectable, and cannot be disabled inadvertently or otherwise.
- FlashArrays' high availability helps satisfy GDPR requirements for keeping subjects' data available. In conjunction with redundant storage networks and application clusters, it can provide end-to-end high availability for personal data in digital form.

Compliance with the digital provisions of GDPR is necessarily an integration of application, server, network, and storage data protection facilities, together with data processor policies for handling and protecting personal data while it is in digital form.

APPENDIX
FLASHARRAY AND GDPR COMPLIANCE

Table 1 lists excerpts from the GDPR articles that relate to digital processing, storage, and transmission of personal data and describes the FlashArray capabilities that help users comply with them.

Table 1: GDPR Text Relationship to FlashArray Properties and Capabilities

Article	Text Excerpt	Relationship to FlashArray
3.1	This Regulation applies to the processing of personal data in the context of the activities of an establishment of a controller or a processor in the Union, regardless of whether the processing takes place in the Union or not.	FlashArrays have almost no optional features. The arrays' high availability, data security, and access controls are the same throughout the product line. Thus, no matter where a FlashArray is deployed, its role in GDPR compliance is as described in the body of this brief.
3.3	This Regulation applies to the processing of personal data by a controller not established in the Union, but in a place where Member State law applies by virtue of public international law.	
4(2)	'processing' means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction;	FlashArrays perform or participate in the highlighted operations.
4(12)	'personal data breach' means a breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data transmitted, stored or otherwise processed;	The body of the brief describes how FlashArrays' always-on data and metadata encryption, in conjunction with network encryption facilities for data in transit minimizes the possibility of data breaches. In addition, credential-based, role-oriented administrative access, together with rigorous audit logging integrated with centralized syslog servers minimizes the possibility of unauthorized processing or misappropriation. Finally, the arrays' built-in protections against single and double component failure represent state of the art protection against accidental loss, destruction, or damage of personal data.
5.1(f)	processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures ('integrity and confidentiality').	
6.4(e)	...the controller shall, in order to ascertain whether processing for another purpose is compatible with the purpose for which the personal data are initially collected, take into account, inter alia: ... the existence of appropriate safeguards, which may include encryption or pseudonymisation.	

Article	Text Excerpt	Relationship to FlashArray
17.2	Where the controller has made the personal data public and is obliged pursuant to paragraph 1 to erase the personal data, the controller, taking account of available technology and the cost of implementation, shall take reasonable steps, including technical measures, to inform controllers which are processing the personal data that the data subject has requested the erasure by such controllers of any links to, or copy or replication of, those personal data.	From a FlashArray standpoint, this provision primarily applies to bulk erasure of entire data sets. Array administrators can bulk-eradicate entire volumes of data, however, application servers must command arrays to erase ranges of data blocks.
24.1	Taking into account the nature, scope, context and purposes of processing as well as the risks of varying likelihood and severity for the rights and freedoms of natural persons, the controller shall implement appropriate technical and organisational measures to ensure and to be able to demonstrate that processing is performed in accordance with this Regulation. Those measures shall be reviewed and updated where necessary.	FlashArray encryption uses the AES-256 algorithm to secure staged and stored data. The arrays' high availability and administrative access control and auditing features help processors comply with the data availability, security, and access control provisions of the regulation.
24.2	Where proportionate in relation to processing activities, the measures referred to in paragraph 1 shall include the implementation of appropriate data protection policies by the controller.	
25.1	Taking into account the state of the art, the cost of implementation and the nature, scope, context and purposes of processing as well as the risks of varying likelihood and severity for rights and freedoms of natural persons posed by the processing, the controller shall, both at the time of the determination of the means for processing and at the time of the processing itself, implement appropriate technical and organisational measures, such as pseudonymisation, which are designed to implement data-protection principles, such as data minimisation, in an effective manner and to integrate the necessary safeguards into the processing in order to meet the requirements of this Regulation and protect the rights of data subjects.	FlashArrays use the AES-256 algorithm to encrypt stored data. The data security, high availability, administrative access control, and auditing features are all standard in every FlashArray. Taken together, they help processors "design in" protection for personal data that complies with the high availability, security, and access control provisions of the GDPR as they implement new processing systems.
29	The processor and any person acting under the authority of the controller or of the processor, who has access to personal data, shall not process those data except on instructions from the controller, unless required to do so by Union or Member State law.	
30.1	Each controller and, where applicable, the controller's representative, shall maintain a record of processing activities under its responsibility. ...several processing activities called out...	FlashArray logs record all administrative actions, including volume creations, resizings, and destructions, host connections and disconnections, and snapshot and replication schedule creations. As such, they provide the informational basis for compliance with this article, especially when integrated with data processors' syslog servers.

Article	Text Excerpt	Relationship to FlashArray
32.1	Taking into account the state of the art, the costs of implementation and the nature, scope, context and purposes of processing as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons, the controller and the processor shall implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk, including inter alia as appropriate: (a) the pseudonymisation and encryption of personal data; (b) the ability to ensure the ongoing confidentiality, integrity, availability and resilience of processing systems and services; (c) the ability to restore the availability and access to personal data in a timely manner in the event of a physical or technical incident; (d) a process for regularly testing, assessing and evaluating the effectiveness of technical and organisational measures for ensuring the security of the processing.	FlashArray encryption of stored data, high availability, administrative access control, and auditing features help processors comply with the data availability, security, and access control provisions of the regulation. Full compliance would also entail network encryption for data in transit, application trustworthiness (e.g., with regard to pseudonymisation), and rigorous policies for data access and handling by agents of the data processor.
34.3(a)	The communication to the data subject referred to in paragraph 1 shall not be required if any of the following conditions are met: (a) the controller has implemented appropriate technical and organisational protection measures, and those measures were applied to the personal data affected by the personal data breach, in particular those that render the personal data unintelligible to any person who is not authorised to access it, such as encryption;	Paragraph 1 of this provision defines the data controller's obligation to notify data subjects of personal data breaches "without undue delay." FlashArray encryption and access control are regarded as "appropriate technical protection measures," so this provision is expected to apply where FlashArrays are used for data storage.
35.1	Where a type of processing in particular using new technologies, and taking into account the nature, scope, context and purposes of the processing, is likely to result in a high risk to the rights and freedoms of natural persons, the controller shall, prior to the processing, carry out an assessment of the impact of the envisaged processing operations on the protection of personal data. A single assessment may address a set of similar processing operations that present similar high risks.	Through its system engineering organization, Pure makes available a series of technical reports and briefs that controllers can use to help assess the impact of proposed operations on personal data protection. A complete assessment, however, would take into account all facets of a proposed operation, including data acquisition, processing, storage, transmission, and eventual destruction.